

Informed Heuristics For Guiding Stem-And-Cycle Ejection Chains

Daniel Harabor and Philip Kilby

NICTA and The Australian National University
7 London Circuit, Canberra, 2601, Australia
{daniel.harabor, philip.kilby}@nicta.com.au

Abstract. The state of the art in local search for the Traveling Salesman Problem is dominated by ejection chain methods utilising the Stem-and-Cycle reference structure. Though effective such algorithms employ very little information in their successor selection strategy, typically seeking only to minimise the cost of a move. We propose an alternative approach inspired from the AI literature and show how an admissible heuristic can be used to guide successor selection. We undertake an empirical analysis and demonstrate that this technique often produces better results than less informed strategies albeit at the cost of running in higher polynomial time.

1 Introduction

Simple to state yet difficult to solve, the Traveling Salesman Problem (TSP) is one of the oldest and most widely studied in the field of combinatorial optimisation. The canonical (or symmetric) TSP requires finding a minimum cost tour of all nodes in a graph $G = (V, E)$ where $V = \{v_1, \dots, v_n\}$ is the set of all nodes and $E = \{(v_i, v_j) | v_i, v_j \in V, i \neq j\}$ is a set of edges connecting them such that cost (or weight) of each edge $c_{ij} = c_{ji}$ and is non-negative for all $(v_i, v_j) \in E$.

Despite receiving much attention the TSP has resisted attempts at developing a polynomial time algorithm to solve it exactly; a situation which is unlikely to change unless $P = NP$ [7]. In response many researchers began focusing their attention on local search methods in order to produce approximate solutions in low polynomial time. By far the most successful of these efforts is the seminal Lin-Kernighan algorithm [8], variations on which have dominated the area for almost three decades. Recently however it has been shown that alternative techniques, using the Subpath Ejection Chain (or SEC) procedure, are able to compete with and in many cases outperform Lin-Kernighan [11,4,2].

SEC proceeds by identifying the *Stem-and-Cycle reference structure*; a spanning subgraph of G which is similar to but not quite a tour (yet from which a tour can always be immediately generated). Using special transition rules one configuration of the reference structure is transformed into another such that a compound neighbourhood of reachable tours is created and from this neighbourhood a cost improving tour is selected. The power of the technique can

be explained, in part, by a connectivity result due to Glover [3] which shows that the neighbourhood generated using Stem-and-Cycle is a super-set of that generated by Lin-Kernighan.

At the heart of the SEC procedure is a simple policy which extends the ejection chain by always choosing to expand the neighbouring configuration with minimum transition cost [3]. While effective it has been suggested that the nearest-neighbour strategy fails to exploit the geometric structure of the TSP and tends to perform poorly on problem instances where nodes appear in clusters [4]. In the AI literature such difficulties are often overcome by extracting knowledge from the problem domain. Known as *informed heuristics* these methods estimate the cost of all remaining steps required to solve the problem. This is in stark contrast to the nearest-neighbour strategy which only considers one step at a time.

We develop an alternative move operator which uses an informed heuristic to rank neighbouring configurations according to their “distance” from an apriori unknown target tour. We compute a 1-tree lower-bound [6] on the cost of the target tour and work on the thesis that this heuristic will guide our algorithm to neighbourhoods which contain lower cost tours and in relatively fewer steps than would otherwise be reachable using SEC.

Our motivation is to show that on many smaller TSP instances ejection chains guided by informed heuristics are able to achieve better results than alternatives based on the nearest-neighbour strategy. Smaller TSPs are of interest because they often appear in many other areas of combinatorial optimisation. For example, in the Vehicle Routing Problem the number of requests which may be assigned to any one route is necessarily small (usually less than 100) due to physical limits on the capacity and range of each truck.

2 Subpath Ejection Chains using Stem-and-Cycle

The Stem-And-Cycle reference structure, as shown in Figure 1(a), is a spanning subgraph of G that consists of a set of nodes $ST = \{v_t, \dots, v_r\}$ forming a simple path called the *stem* which is connected to another set of nodes $CY = \{v_r, v_{s_1}, \dots, v_{s_2}, v_r\}$ termed the *cycle*. The node connecting the stem to the cycle, v_r , is called the root while the node at the other end of the stem, v_t , is known as the tip. v_r must always be connected to at least two unique nodes on the cycle v_{s_1} and v_{s_2} which are respectively designated the left and right subroots.

The main idea behind Stem-and-Cycle (whose theoretical underpinnings may be found in [3]) is to search for a cost improving tour by re-arranging subpaths within the reference structure. This is achieved by defining a series of transition rules, or moves, that transform one Stem-and-Cycle configuration into another. Each move adds an edge (v_t, v_i) that connects the tip to a node on the stem or cycle and deletes an adjacent edge (v_i, v_j) such that the resulting spanning graph remains connected. When $v_t = v_r$ the structure is said to be degenerate and forms a valid tour. Further, although the Stem-and-Cycle structure is not itself a feasible solution to the TSP it is always possible to generate two such

tours or *trial solutions*, by adding an edge (v_t, v_{s_1}) or (v_t, v_{s_2}) and deleting the edge (v_{s_1}, v_r) or (v_{s_2}, v_r) as appropriate.

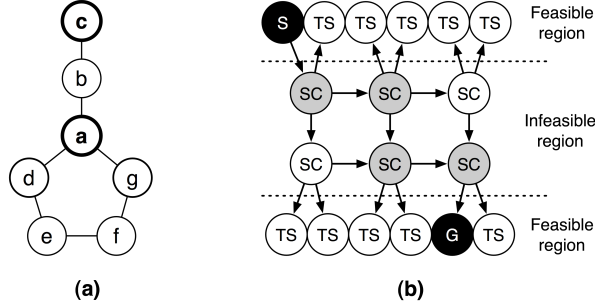


Fig. 1. (a) The Stem-and-Cycle reference structure. Node c is the tip, node a the root and nodes d and g the left and right subroots respectively. **(b)**. SEC begins at tour S and creates a 4-level ejection chain (each level is denoted by a gray node) before choosing G as the best reachable trial solution (TS).

The Stem-and-Cycle reference structure is used as the basis for the Subpath Ejection Chain (or SEC) move operator. We briefly outline it here due to its central importance in developing other ideas in this paper; a full description, including pseudocode, is given in [11].

In broad terms SEC explores a neighbourhood by producing a series of inter-related moves that generate one Stem-and-Cycle configuration from another in a chain-like fashion. Each move constitutes a new level of the ejection chain and at each level it is necessary to evaluate the set of all legitimate Stem-and-Cycle configurations in order to select one which will further extend the chain. When the ejection chain has reached a maximal size the best trial solution reachable from any of its levels is selected as the new start tour. The process is repeated until SEC reaches a local minimum where no cost improving tour can be found. Figure 1(b) serves to illustrate this idea.

In order to guide SEC toward promising regions of the search space two important constraints are applied:

1. Each added edge is marked interchangeably as either black or white. White edges are assumed to belong to a target tour into which we seek to transform the starting tour. Thus, white edges, once added, cannot be subsequently deleted. Black edges are assumed to not belong to the target tour and may be deleted as part of a future move.
2. Each deleted edge is marked “tabu” and cannot appear in a subsequent level of the ejection chain.

The first constraint seeks to focus the search toward an apriori unknown target tour which may or may not be cost improving. It also serves to limit the length of the ejection chain to a maximum of $2n$ levels, where $n = |V|$.

The second constraint prevents the search from cyclical backtracking. Any move which violates one or both constraints is deemed to be illegitimate and not considered.

3 Growing Ejection Chains Using Heuristic Search

The SEC algorithm searches for cost improving tours by creating a series of multi-level ejection chains. At each level of a chain the algorithm generates a set Γ of up to $2n$ legitimate Stem-and-Cycle configurations [3]. Each $x \in \Gamma$ is evaluated using NN , a nearest-neighbour heuristic function modeled as:

$$NN(x) = ejectionValue + \min(trialValue)$$

The first term, *ejectionValue*, refers to the cost of the move generating x from the current Stem-and-Cycle configuration. The second term, $\min(trialValue)$, is the minimum cost associated with generating one of two the trial solutions reachable from x . Since each $NN(x) | x \in \Gamma$ requires constant time to evaluate a maximally sized ejection chain with up to $2n$ levels may be constructed in no more than $O(4n^2)$ time.

Although SEC runs very fast its nearest-neighbour heuristic is not very accurate. In particular, there is a risk that SEC will choose a successor which appears promising but ultimately leads to a section of the search space that contains no cost improving tour. To address this issue we propose an alternative move operator that is based, in part, on a well known algorithm from the heuristic search literature: A* [5]. Developed in the context of the single-source shortest-path problem, A* makes use of an informed heuristic function that evaluates nodes in the search space by computing a lower-bound on the remaining distance to reach a known goal node. In the canonical implementation the final or *f-cost* associated with an arbitrary node x is given by the function:

$$f(x) = g(x) + h(x)$$

where h is the heuristic cost and g is the weight of the path from the start node to x . Provided h is *admissible* (i.e. it never overestimates the distance to the goal) this approach is guaranteed to find a shortest path to the goal if one exists. Applying such an algorithm to solve the TSP is not straightforward however. In particular, it is difficult to measure the “distance” between an arbitrary tour and the optimal tour as the latter is not known apriori. Another difficulty arises from the fact that A* is a systematic tree-search algorithm; when applied to the TSP it exhibits exponential space complexity requirements which make it infeasible for solving all but the most trivial instances [10,9]. In the remainder of this section we will address each of these problems in turn and in the process develop a new Stem-and-Cycle move operator that constructs ejection chains with the help of a an informed heuristic function.

3.1 Problem Formulation

We begin by recalling that every ejection chain created by the Stem-and-Cycle move operator must respect a set of edge exclusion constraints (namely those edges which have been previously deleted and cannot be re-added) and a set of edge inclusion constraints (namely those edges marked as white which cannot be deleted). Combined, these constraints guide the ejection chain toward a particular but apriori unknown tour. We designate this tour g and the associated start tour, from which the algorithm begins, s . Let these tours correspond to A^* 's goal node and start node respectively.

We define our search space as an undirected meta-graph $G_m = (V_m, E_m)$ where V_m represents the set of nodes in the graph and E_m the set of edges. Let each node $x \in V_m$ represent a valid Stem-and-Cycle configuration. We will say that an edge $(x, y) \in E_m$ exists between two nodes, x and y , if the Stem-and-Cycle move operator can induce y from x (or vice versa) in a single move. The weight of each edge is thus equal to the cost of the associated Stem-and-Cycle move that induces one of its endpoints from the other. Under this formulation $s, g \in V_m$ and our goal is to find a path $\Pi(s, g)$ connecting them such that the weight of g is minimised:

$$\Pi(s, g) : \min(g)$$

The existence of the path is guaranteed by a result due to Glover [3] which shows that Stem-and-Cycle can generate a tour from any other using not more than $2n$ moves.

3.2 Modeling The Cost Function

Recall that A^* ranks nodes by computing a final value which is the sum of two cost functions: g and h . We proceed by developing TSP-specific analogues for each of these.

First we consider the problem of calculating a g -cost which requires exactly measuring the distance from s to some candidate successor node $c \in V_m$. To achieve this we can simply take the sum of the weights of all edges designated white appearing in c . We reason that a node which contains more white edges is in some sense "closer" to g than one containing fewer white edges.

Meanwhile, calculating a h -cost is analogous to computing a lower-bound on the cost of all white edges which have yet to be added in order to transform c into g . We accomplish this by taking the weight of all non-white edges in a *constrained 1-tree*.

Conceptually similar to the canonical 1-tree [6] (the construction of which we outline in Figure 2) the constrained 1-tree differs in that it contains all edges from a given include set and none of the edges from a given exclude set. Algorithm 1 describes the steps necessary to create this data structure. Note that this idea was originally suggested by Pohl [10] but not described in any detail. As such we believe this work is the first elaborate study of the constrained 1-tree.




Fig. 2. (a) Begin the 1-tree by first constructing an MST of all tour nodes except one (here, node i). (b) Complete the 1-tree by connecting node i to the rest of the spanning tree via the two smallest weight edges incident with it (here, edges (i, a) and (i, e)).

Our implementation takes a set of nodes $V = \{v_1, \dots, v_n\}$ and proceeds by creating a MST on the subset of nodes $\{v_2, \dots, v_n\}$. However we augment the basic MST construction algorithm of Prim (as described in [1]) such that it respects all edge constraints to which the move operator is subject. We make two changes:

1. In lines 19-21 of Algorithm 1: white edges are added to the heap are given the highest priority. This ensures that all such edges are immediately selected for inclusion into the MST.
2. In lines 23-25 of Algorithm 1: before any non-white edge is added to the heap we make sure it is not on the excluded edges list. If it is, we immediately discard it from further consideration.

We also apply the following rules at line 16 of Algorithm 1 (where we connect the excluded node v_1 to the MST and complete the 1-tree):

1. Any white edge incident with node v_1 which appears in the included edges list is automatically selected, regardless of cost.
2. Any non-white edge incident with node v_1 which appears in the excluded edges list is automatically discarded, regardless of cost.
3. The remaining edges incident with node v_1 are ranked by their weight. From this set we always select the one with smallest weight.

3.3 On The Admissibility Of The Constrained 1-Tree Heuristic

The canonical 1-tree always lower-bounds the optimal tour and thus the the goal configuration g . However, it is not immediate that the constrained 1-tree, as constructed by Algorithm 1, is also a lower-bound on g . In particular we would like to show that the heuristic function h , which is based on the sum cost of non-white edges in a constrained 1-tree, is a lower-bound on the sum cost of any as-yet-unknown set of white edges which appear in the goal tour g . This property is required to ensure the admissibility of the heuristic function h .

Algorithm 1 Constrained 1-Tree**Require:** $G = (V, E)$, $ExcludedEdges$, $IncludedEdges$

```

1:  $v_1 \leftarrow \text{getNodeAtIndex}(V, 0)$ 
2:  $mstroot \leftarrow \text{getNodeAtIndex}(V, 1)$ 
3:  $Neighbours \leftarrow V - \{v_1 \cup mstroot\}$ 
4:  $Heap \leftarrow \text{initHeap}(Neighbours)$ 
5:  $\text{processNeighbours}(mstroot)$ 
6: while  $Neighbours \neq \emptyset$  do
7:    $n \leftarrow \min(Heap)$ 
8:   if  $\text{parent}(n) = \emptyset$  or  $n = \emptyset$  then
9:     return  $\emptyset$ 
10:  if  $\text{makeEdge}(\text{parent}(p), n) \in ExcludedEdges$  then
11:    return  $\emptyset$ 
12:   $\text{parent}(n).addChild(n)$ 
13:   $Neighbours \leftarrow \{Neighbours - n\}$ 
14:   $\text{processNeighbours}(n)$ 
15: return  $\text{make1Tree}(mstroot, v_1)$ 

  procedure  $\text{processNeighbours}(n)$ 
16: for all  $x \in Neighbours$  do
17:    $e \leftarrow \text{makeEdge}(n, x)$ 
18:   if  $e \in IncludedEdges$  then
19:      $Heap.priority(x, 0)$ 
20:      $\text{parent}(x) \leftarrow n$ 
21:   else
22:     if  $e \notin ExcludedEdges$  and  $\text{priority}(x) < \text{cost}(e)$  then
23:        $Heap.priority(x, \text{cost}(e))$ 
24:        $\text{parent}(x) \leftarrow n$ 

```

Lemma 1. *For all $c \in V_m$ it is always possible to construct a constrained 1-tree.*

Proof. To prove this claim it is sufficient to observe that every Stem-and-Cycle configuration, including the optimal tour, is itself a 1-tree. \square

Lemma 2. *The constrained minimum spanning tree associated with an arbitrary node $c \in V_m$ is a lower-bound on the goal configuration g .*

Proof. We proceed with a structural proof by induction on the spanning tree constructed for c . First we will analyse two distinct possibilities for the base case: $c = s$ or $c = g$. We then proceed with an analysis of the general case in which c is subject to an arbitrary number of edge inclusion and edge exclusion constraints. In each case we will show that Algorithm 1 produces a spanning tree which is a lower-bound on g .

Base Case 1: $c = s$. In this case there are no edge inclusion or exclusion constraints. Algorithm 1 mirrors the behaviour of Prim's algorithm and thus creates an unconstrained spanning tree of minimum weight. Clearly such a structure is a lower-bound on g .

Base Case 2: $c = g$. In this case every edge in c is white and appears in the set of edges that must be included. Algorithm 1 constructs the constrained MST by looking for minimum cost edges that connect a leaf of the spanning tree to some other node that is not yet in the tree. If any edge from this candidate set is white it is assigned the highest priority which guarantees it will be selected the next time the spanning tree is extended. Since every node in c is incident with at least one white edge it follows that each node is added to the spanning tree by way of a white edge. As the constrained MST contains only white edges its weight must be a lower-bound on g .

Inductive Case: c contains some number of white and non-white edges. We know from Base Case 2 that every white edge appearing in g also appears in the constrained MST associated with c . It remains to show that the set of non-white edges appearing in the constrained MST is minimum.

When considering a non-white edge Algorithm 1 will rank candidates that have not been explicitly excluded by their weight and always select for inclusion into the spanning tree the edge with smallest weight. Consequently, no edges which have been deleted can appear in the spanning tree and of the remaining candidate edges only those edges of least cost which do not induce a cycle are selected when extending the tree. The set of non-white edges included in the spanning tree is thus minimum and must be a lower-bound on the weight of any white edges which later appear in g . Thus the weight of the constrained MST is a lower-bound on g .

□

Theorem 1. *The weight of a constrained 1-tree associated with some arbitrary node $c \in V_m$ is a lower-bound on the weight of the goal configuration g .*

Proof. We recall that a constrained 1-tree on the set of nodes $\{v_1, \dots, v_n\} \in c$ is a two-step process requiring the union of a constrained MST on the nodes $\{v_2, \dots, v_n\}$ with the two least-cost (possibly white) valid edges incident with v_1 . We know from Lemma 2 that the weight of the constrained MST is a lower-bound on g . It remains to show that the two edges connecting v_1 to the spanning tree are a lower-bound on the cost of two edges incident with v_1 in the goal configuration g .

This operation follows the edge selection policy outlined in Section 3.2; i.e. we preference white edges above all others, regardless of cost, disconsider excluded edges and rank the remaining edges by weight. Thus the two edges selected to connect v_1 to the constrained MST must be a lower-bound on the cost of the two edges incident with v_1 in g . From this we conclude that the constrained 1-tree constructed by Algorithm 1 is itself a lower-bound on g .

□

Theorem 1 is sufficient to show that a h is admissible. Furthermore, each generated node evaluated by h will be subject to a larger set of edge constraints than its parent. From this it is readily shown that h is also *monotonic* and thus guaranteed to be non-increasing as the search approaches the target tour g .

3.4 The Informed Subpath Ejection Chain Algorithm

We have combined the Stem-and-Cycle reference structure with the constrained 1-tree lower-bound in order to develop a TSP-specific analogue of A^* . If we apply the canonical form of the algorithm to G_m , always expanding the node with the least lower-bound, we are guaranteed to find the optimal tour. However, this operation is likely to be computationally intractable unless our heuristic function h is able to compute a near-perfect estimate of the remaining distance to reach the optimal tour. If this is not the case A^* will expand an exponential number of nodes [9].

Since we are interested in an approach that is guaranteed to run in low polynomial time we modify the canonical A^* algorithm by reducing the branching factor of each node to 1. This makes our new algorithm analogous to SEC [3,11] however instead of always choosing the neighbour with minimum transition cost as the next level of the ejection chain we instead choose the neighbour which minimises the cost of our evaluation function f . We term the resulting technique Informed SEC (or ISEC for short) and outline it in Algorithm 2.

Algorithm 2 Informed Subpath Ejection Chain Algorithm (ISEC)

Require: s

```

1:  $next \leftarrow s$ 
2:  $bestTour \leftarrow \text{minTrialSolution}(s)$ 
3: while  $next \neq \emptyset$  do
4:   if  $\text{minTrialSolution}(next) < bestTour$  then
5:      $bestTour \leftarrow \text{minTrialSolution}(next)$ ;
6:    $\Gamma_{next} \leftarrow \text{expand}(next)$ ;
7:    $next \leftarrow \emptyset$ 
8:   for all  $x \in \Gamma_{next}$  do
9:     if  $next = \emptyset$  then
10:       $next \leftarrow x$ 
11:     else
12:       if  $f(x) < f(next)$  then
13:          $next \leftarrow x$ 
14: return  $bestTour$ 

```

4 Complexity

We analyse the running time of the ISEC algorithm in terms of node expansions. When a node $x \in G_m$ is expanded we generate the set Γ_x which contains all legitimate Stem-and-Cycle configurations reachable in a single move. We then select $\min(\Gamma_x)$ as the next node to be expanded and the process continues until a maximal size ejection chain has been built. Since every second move adds to x a white edge which cannot be deleted the maximum number of nodes that ISEC can expand per ejection chain is $2n$. Additionally, the size of Γ_x is itself

not more than $2n$ which gives us an $O(4n^2)$ bound on the total number of nodes that must be generated per ejection chain.

For each node $y \in \Gamma_x$ we call Algorithm 1 to create a constrained 1-tree. Using a derivative of Prim's MST algorithm we first compute a constrained MST on the nodes $\{v_2, \dots, v_n\} \in V$. This is an $O((n-1)^2 \times \log_2(n-1))$ operation on a complete graph such as the Symmetric TSP. To finish the 1-tree we connect node $v_1 \in V$ to the constrained MST via its two least-cost edges. However, we have to exclude from consideration any edges which have been previously deleted and prioritise any edges which have been marked white. Finding two edges which satisfy these constraints is an $O(n-1)$ operation, even if we pre-sort the set of edges incident with v_1 . Thus the total time required to compute a constrained 1-tree is $O((n-1)^2 \times \log_2(n-1) + (n-1))$.

If we assume that each ejection chain constructed by ISEC will require a maximal number of node expansions then the total running time of the algorithm is $O(4n^2 \times [(n-1)^2 \times \log_2(n-1) + (n-1)])$. We may roughly approximate this as $O(4n^4 \times \log_2(n) + 4n^3)$.

Clearly the performance of ISEC is limited by the time required to compute a constrained 1-tree. However, it is not strictly necessary to compute a new 1-tree for each generated node. Consider the following argument:

Let $p, c \in V_m$ be two Stem-and-Cycle configurations such that c is generated as a result of expanding p . Such a move adds edge e_a and deletes the edge e_d . Next, let τ_p and τ_c be the associated constrained 1-trees which are constructed to evaluate p and c respectively. We proceed with a discussion on the effect that the Stem-and-Cycle move generating c from p has on the structure of τ_c with respect to τ_p . In particular, we identify four distinct types of changes which may occur:

Change 1: $e_a \notin \tau_p$ and $e_d \notin \tau_p$. In this case two possibilities exist, depending on whether e_a is white or not.

- i) If we suppose e_a is not white then we are forced to include in τ_c any edges which were not already in τ_p . Further, since $e_d \notin \tau_p$ we do not have prevent any edges appearing in τ_c which were present in τ_p . Thus, $\tau_p = \tau_c$ and we do not have to generate a separate 1-tree for c .
- ii) If we suppose e_a is white then we are forced to include it in τ_c . However as $e_a \notin \tau_p$ we cannot re-use τ_p in place of τ_c and still retain the admissibility property of the heuristic function h . Thus we must generate τ_c from scratch.

Change 2: $e_a \notin \tau_p$ and $e_d \in \tau_p$. In this case the deletion of e_d divides τ_p into two separate subtrees which must be re-connected with a least-cost edge that has not been previously deleted. Finding such an edge is an operation requiring $O(n^2)$ time and thus asymptotically equivalent with computing t_c from scratch. The status of e_a (black or white) is of no consequence.

Change 3: $e_a \in \tau_p$ and $e_d \notin \tau_p$. In this case two possibilities exist, depending on whether e_a is white or not.

- i) If e_a is a white edge its inclusion into τ_c is forced. Consequently, Algorithm 1 will produce a 1-tree such that $cost(\tau_c) \geq cost(\tau_p)$. However, as $e_a \in \tau_p$ it must be that τ_p respects the same set of constraints to which τ_c is subject. Thus, τ_p is also a valid lower-bound for c and we may use it in favour of generating τ_c .
- ii) If e_a is a black edge the same argument as before holds; τ_p is a direct lower-bound for c and we may use it in favour of generating τ_c .

Change 4: $e_a \in \tau_p$ and $e_d \in \tau_p$. The same argument which we used for Change 2 is also applicable here. The deletion of e_d divides τ_p into two separate subtrees which must be re-connected with a least cost-edge. This operation is asymptotically equivalent to generating τ_c from scratch. The status of e_a (black or white) is of no consequence.

These results are interesting because they show that in half of all node generation scenarios it is not necessary to compute a 1-tree. In spite of this the theoretical running time of ISEC remains unchanged and in the worst case we still need to compute a constrained 1-tree for every node generated. We have observed that in practice this scenario is unlikely; as our experimental results will show a speedup of 50% is not uncommon. There are however some theoretical side-effects which merit discussion. In particular, notice that in Change 3i although τ_c can be shown to obey the same set of constraints as τ_p its cost may in fact be higher. In such cases the accuracy of the heuristic function h is diminished as not all constraints were taken into consideration when computing $h(c)$. We explore this issue empirically and show that in practice it is not a significant drawback.

5 Experimental Setup

We implemented and tested SEC (as described in [11]), ISEC and fast-ISEC (a variant of ISEC using the speedup technique described in Section 4) on 30 small-to-medium problems from TSPLIB [12]. These ranged in size from 29 to 200 cities. We generate 10 experiments per problem by randomly perturbing the initial order of cities. We evaluate the performance of both algorithms on each of these 10 instances making for a total of 600 distinct experiments. All tests were undertaken on an Intel Core2Duo machine with 2GB of RAM running OSX 10.5.7.

Our main objective is to compare the relative quality of successor nodes chosen during a neighbourhood search. To this end we ran a single iteration of both SEC and ISEC and measured the length of the best corresponding trial solution. Note that in the case of ISEC an iteration refers to the construction of a single ejection chain of maximal length. In the case of SEC an iteration may consist of many such ejection chains; the algorithm only terminates when a local minima is reached.

6 Results

Table 1 summarises the quality of computed solutions for the SEC, ISEC and fast-ISEC algorithms. We evaluate algorithmic performance by taking the length of each solution found and measuring the relative percentage deviation from the optimal solution. In each case we give figures for the best and average solutions found (columns Min and Mean respectively). We also provide average running times for each algorithm (in columns CPU) which are taken as wall-clock time and measured in seconds. Finally we evaluate the average quality of the starting tour (column STQ) as measured in multiples of the optimal solution length.

Table 1. Quality of solutions (% above optimal) and CPU time (in seconds) for SEC, ISEC and fast-ISEC.

Problem	Optimal	STQ	ISEC			fISEC			SEC		
			Min	Mean	CPU	Min	Mean	CPU	Min	Mean	CPU
bayg29	1610	2.96	14.60	23.34	0.13	14.60	24.87	0.09	5.59	15.11	0.00
berlin52	7542	400.20	19.49	30.83	0.98	18.86	37.23	0.58	17.32	27.61	0.01
bier127	118282	521.52	28.58	35.14	23.97	27.50	35.92	12.88	27.68	35.88	0.09
ch130	6110	764.88	27.72	39.19	26.91	27.79	40.55	14.53	42.51	58.72	0.08
ch150	6528	811.33	21.71	43.09	46.57	32.10	50.97	24.51	62.40	82.20	0.11
d198	15780	1202.76	29.38	39.43	135.78	29.38	39.35	70.88	60.32	67.72	0.17
eil101	629	543.48	17.20	31.63	10.30	26.24	35.56	5.71	25.90	35.10	0.05
eil51	426	378.29	11.69	24.36	0.91	11.25	29.46	0.54	11.03	22.09	0.01
eil76	538	459.30	13.61	32.18	3.74	20.28	36.99	2.07	26.47	41.20	0.02
kroA100	21282	786.14	25.01	43.47	10.19	38.45	49.29	5.54	32.58	50.81	0.04
kroA150	26524	957.21	29.64	41.96	47.45	34.42	46.59	25.44	45.29	56.03	0.14
kroA200	29368	1165.43	38.66	47.27	150.71	38.83	50.28	77.80	53.40	81.68	0.18
kroB100	22141	754.03	22.62	35.92	10.26	22.71	39.09	5.63	37.71	57.51	0.05
kroB150	26130	1002.18	31.70	44.77	47.25	31.70	49.86	24.71	58.01	66.74	0.12
kroB200	29437	1140.09	45.67	54.03	145.55	38.80	49.75	77.06	64.23	83.93	0.20
kroC100	20749	823.89	31.58	42.36	10.36	18.79	44.53	5.68	38.71	55.75	0.04
kroD100	21294	776.11	31.97	39.82	10.49	33.29	48.22	5.59	37.33	51.13	0.04
kroE100	22068	756.53	29.91	46.35	10.44	27.99	46.89	5.59	45.20	73.82	0.04
lin105	14379	845.14	42.62	58.56	11.85	37.15	63.15	6.57	51.57	87.18	0.04
pr107	44303	1297.00	33.95	64.24	12.08	34.10	58.51	6.59	38.20	49.27	0.05
pr124	59030	1176.19	37.37	61.71	21.82	45.05	68.07	11.70	64.38	103.39	0.06
pr136	96772	832.38	27.40	42.08	30.18	37.80	47.77	16.42	49.36	68.92	0.07
pr144	58537	1361.50	40.13	74.17	37.10	39.21	66.39	20.24	65.06	90.74	0.09
pr152	73682	1398.02	49.02	63.52	45.64	35.76	59.54	25.10	48.74	70.63	0.13
pr76	108159	522.91	24.65	34.35	3.45	28.65	41.75	2.02	36.74	47.90	0.03
rat195	2323	968.21	33.77	41.01	126.33	35.69	42.24	66.77	80.39	95.76	0.19
rat99	1211	682.64	24.34	40.77	9.32	24.34	43.57	5.14	39.38	74.09	0.04
rd100	7910	709.22	24.36	39.70	10.45	26.62	39.76	5.65	38.23	57.64	0.04
st70	675	560.14	22.69	34.73	2.83	26.36	40.15	1.61	33.60	54.91	0.02
u159	42080	1059.40	46.24	54.21	53.80	39.20	54.18	28.59	72.48	93.99	0.12
Mean		821.97	29.24	43.47	35.23	30.10	46.02	18.71	43.66	61.91	0.08

Looking at the Table 1 we see that with respect to the quality of the average solution found ISEC outperforms SEC in 26 of the 30 problem instances. In particular, the tours found by ISEC are on average 18.44% closer to optimal. A similar observation holds if we look at the quality of minimum solutions found. As before, ISEC’s performance dominates SEC – this time in 25 of the 30 problem instances – by an average of 14.4%. Similar results are seen in the case of fISEC however the average and minimum solution quality improvements over SEC are slightly lower at 15.89% and 13.56% respectively. As this is quite a coarse analysis we turn our attention to Figure 3(a) which shows the spread of results for mean solution quality.




Fig. 3. (a) Variability of mean solution quality (measured as % above optimal). (b) The effect of start tour quality on solution quality.

Here we notice that the inter-quartile range of tours found by ISEC is just 11% and includes no solutions more than 46% from optimal. Meanwhile SEC exhibits quite different behaviour; the inter-quartile range encompasses a very wide 39% range and includes tours which are up to 80% above optimal. To explore this effect we turn our attention to Figure 3(b) in which we measure the effect of start tour quality on solution quality.

Using a simple least squares regression analysis we are able to see that the performance of both algorithms appears to be strongly correlated to the quality of the starting tour. In particular, notice that as the “distance” between the optimal tour and the start tour increases there is a corresponding drop in the mean quality of the best solution found. This suggests that such problems are “harder” than others, requiring the ejection of a greater number of edges before a good quality solution is found. As before, ISEC performs much better than SEC almost across the board. This is directly attributable to our informed heuristic function which, it appears, is able to find promising regions of the search space much more consistently than SEC’s nearest-neighbour method.

Next, we analyse the search effort of SEC and ISEC which we measure in terms of nodes expanded and nodes generated. Looking at Figure 4 we notice that in the case of both metrics the performance of the two algorithms scales linearly with problem size. In the case of ISEC this is because the algorithm only constructs a single ejection chain; thus no more than $2n$ nodes can ever be expanded. By comparison, SEC tends to expand (and consequently generate) an order of magnitude more nodes. From this we may infer that on average approximately 10 ejections chains must be constructed before the algorithm reaches a local minima and terminates.




Fig. 4. Total search effort (SEC vs. ISEC).

An interesting dichotomy appears when we contrast Figure 4 with the corresponding CPU times from Table 1. We see that although ISEC generates much fewer nodes than SEC it is up to 2 orders of magnitude slower. We may attribute this behaviour to ISEC’s evaluation function which must compute a 1-tree for every node generated. This effectively limits ISEC to solving much smaller instances than SEC. Our goal however is not to be able to compute solutions faster than SEC but rather to illustrate the power of informed heuristics for guiding ejection chain algorithms. To that end we believe ISEC’s ability to find alternative neighbourhoods which contain demonstrably higher quality tours makes it a promising alternative. Even so, we can report that the running time of all three algorithms can be improved by artificially limiting the maximum length of each ejection chain from $2n$ to n . This change has a minimal effect on solution quality (the same trends were observed) but halves the CPU times in Table 1.

7 Conclusion

We have shown that it is possible to generate Stem-and-Cycle ejection chains using an informed successor selection strategy. We develop the constrained 1-tree heuristic which computes a lower-bound on an apriori unknown goal tour

with respect to a set of edge inclusion and edge exclusion constraints. We show that this approach is both admissible and monotonic, never overestimating the cost to reach the goal tour. We combine the constrained 1-tree heuristic with the Subpath Ejection Chain (SEC) procedure to develop a new informed move operator which we term ISEC. We show that ISEC, using only a single ejection chain, consistently computes higher quality solutions to the TSP than SEC which generates many more ejection chains. We give complexity results for ISEC and show that although its worst-case theoretical running time is $O(4n^4 \times \log_2(n) + 4n^3)$ we are able to use it as a feasible alternative to SEC for TSP instances of up to 100 nodes.

Future work involves applying ISEC to similar size TSPs featuring precedence and time window constraints. Often appearing as subproblems in other areas of combinatorial optimisation (for example the Vehicle Routing Problem with Time Windows), these domains are interesting because the set of feasible solutions that may be found is much smaller than in the canonical TSP.

7.1 Acknowledgements

We would like to thank Adi Botea for his insightful comments and assistance during the development of this work. NICTA is funded by the Australian Government as represented by the Department of Broadband, Communications and the Digital Economy and the Australian Research Council through the ICT Centre of Excellence program.

References

1. Thomas H Cormen, Charles E Leiserson, Ronald L Rivest, and Clifford Stein. *Introduction to Algorithms*. MIT Press, Cambridge, MA, second edition, 2001.
2. D. Gamboa, C. Rego, and F. Glover. Implementation analysis of efficient algorithms for the traveling salesman problem. *Computers and Operations Research*, (33):1154–1172, 2006.
3. Fred Glover. Ejection chains, reference structures and alternating path methods for traveling salesman problems. *Discrete Applied Mathematics*, 65:223–253, 1996.
4. Gregory Gutin, Abraham Punnen, Alexander Barvinok, Edward Kh. Gimadi, and Anatoliy I. Serdyukov. *The Traveling Salesman Problem and Its Variations*. Kluwer Academic Publishers, 2002.
5. P E Hart, N J Nilsson, and B Raphael. A formal basis for the heuristic determination of minimum cost paths. *IEEE Transactions on Systems Science and Cybernetics*, (4):100–107, 1968.
6. M. Held and R. Karp. The traveling salesman problem and minimum spanning trees. In *Operations Research*, volume 18, pages 1138–1162. 1970.
7. M.R. Garey & D.S. Johnson. *Computers and intractability: A guide to the theory of NP-completeness*. Freeman, San Francisco, 1979.
8. S. Lin and B. W. Kernighan. An effective heuristic algorithm for the traveling salesman problem. *Operations Research*, 21:498–516, 1973.
9. Judea Pearl. *Heuristics: intelligent search strategies for computer problem solving*. Addison-Wesley, Boston, MA, USA, 1984.

10. I. Pohl. Practical and theoretical considerations in heuristic search algorithms. In E. W. Elcock and D. Michie, editors, *Machine Intelligence 8*. Ellis Horwood Ltd and John Wiley & Sons, 1977.
11. C. Rego. Relaxed tours and path ejections for the traveling salesman problem. *European Journal of Operations Research*, (106):522–538, 1998.
12. Gerhard Reinelt. TSPLIB—A Traveling Salesman Problem Library. *INFORMS Journal on Computing*, 3(4):376–384, 1991.